CRT-2000-23

Centre de recherche sur les transports (C.R.T.)
Centre for Research on Transportation

A TABU SEARCH ALGORITHM FOR SELF-HEALING RING NETWORK DESIGN

by

Bernard Fortz
Patrick Soriano
Christelle Wynants

May 2000

101-01794/

CRT-2000-23

*A TABU SEARCH ALGORITHM FOR SELF-HEALING RING NETWORK DESIGN**

by

Bernard Fortz[1]
Patrick Soriano[2]
Christelle Wynants[3]

* *This research was supported in part by the Fonds pour la formation de chercheurs et l'aide à la recherche (FCAR), and Faculty of Management, McGill University.*

[1] Department of Quantitative Economics, Faculty of Economics and Business Administration, Maastricht University, THE NETHERLANDS and Service de Mathématiques de la Gestion, Institut de Statistique et de Recherche Opérationnelle, Université Libre de Bruxelles, Brussels, BELGIUM <bfortz@smg.ulb.ac.be>

[2] Service des Méthodes quantitatives de gestion, École des Hautes Études Commerciales and Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, CANADA H3C 3J7 <patrick@crt.umontreal.ca>

[3] Service de Statistique et Méthodes Quantitatives de Gestion, Institut de Mathématiques, Université de Liège, Liège, BELGIUM <cwynants@ulg.ac.be>

Dépôt légal – Bibliothèque nationale du Québec,
Bibliothèque nationale du Canada, 2000

Abstract

We consider the problem of designing self-healing rings in order to protect the transmission of telecommunication demands in a zonal network. This problem stems from a real application with operational constraints such as dual homing and hop limit per ring. A modeling approach taking into account ring interactions is proposed as well as a tabu search heuristic for solving it. Computational results for a comprehensive set of real and randomly generated instances are presented.

Keywords: telecommunication, network design, survivability, self-healing rings, tabu search

Résumé

Dans cet article nous traitons le problème de conception d'anneaux auto-cicatrisants (*self-healing rings* en anglais) afin de protéger la transmission de signaux dans un réseau de télécommunications zonal. Ce probléme provient d'un contexte d'application partique comportant diverses contraintes opérationnelles telles que l'inclusion de deux points d'interconnection par anneau (*dual homing*) et une limite maximale sur le nombre de points de demande pouvant être connectés à un même anneau (*hop limit*). Une approche de modélisation tenant compte des interactions entre les différents anneaux est proposée de même qu'une méthode de recherche avec tabou pour résoudre ce modéle. La performance de cette méthode est ensuite évaluée sur la base d'un large ensemble d'instances comportant des données réelles et d'autres générése aléatoirement.

Mots-clés: télécommunications, conception de réseaux, résistance aux pannes, anneaux auto-cicatrisants, recherche avec tabous

1 Introduction

A network is said to be *survivable* if traffic interrupted by the failure of some of its elements can be rerouted via spare or excess capacity specifically placed in the network for that purpose. In addition, if that rerouting process can be automated and the network can reconfigure itself to cope with failures without human intervention, then the network is said to be *self-healing*. Of course, designing a network to survive any type of failures including those involving several elements at the same time would result in designs of prohibitive cost. However, since it is generally considered that failures affecting more than one element at a time are extremely improbable, operators instead define a restricted set of realistic failure scenarios for which the network needs to be survivable. This set usually includes all single link and node failures and, possibly some multi-link or node failure scenarios considered critical. The design process then proceeds under the fundamental hypothesis that only one failure needs to be handled at any given time and produces a network that has the topology and spare capacity required to route all the affected traffic "around" these failures (see Cosares *et al.* [6]).

With the advent of SDH/SONET networks and of Add-Drop Multiplexers (ADMs), a new protection technique, known as *Self-Healing Ring* (SHR), was introduced which maintains very fast reconfiguration times while achieving lower redundant capacity requirements (see Sosnosky and Wu [28], Wu, Kolar and Cardwell [34]). Demand nodes are grouped together forming a closed loop or cycle in the network. They are then connected by using two (eventually more) separate fibers, each carrying signals in opposite directions. At each node connected to the ring, an ADM having an appropriate transmission speed or capacity is installed allowing signals to access or leave the ring. Hence, within such a ring architecture, there always exist two link and node disjoint paths connecting any pair of nodes belonging to the ring. All traffic flowing through a ring is therefore protected against any single failure (as well as some multiple failures) of the links or nodes forming it by providing enough spare capacity on the alternate path of each demand. When a failure is detected within a ring, the ADMs on the ring react and switch the affected traffic to the protection path.

There are in fact several types of SHR depending on the way the reconfiguration of the traffic is performed and which type of protocol is used. Note that all ADMs on a given ring have identical capacities and this is referred to as the ring *capacity* or size. This capacity must be sufficient to accommodate the maximum edge flow among the edges belonging to the ring. The amount of spare capacity required is then simply equal to the ring *working* capacity. However for determining the appropriate ring capacity for a given set of demands, SHRs are classified as belonging to one of two categories depending on the way the working

traffic may be routed on them under normal conditions (see Baudron, Khadr and Kocsis [4], Wu [35]).

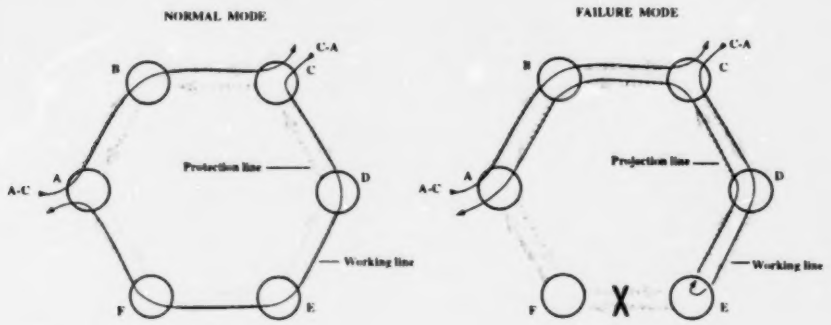

Figure 1: Unidirectional or Dedicated Protection SHR.

In the first type, called *unidirectional* SHR, all traffic is routed in the same direction along one fiber called the working fiber (e.g. clockwise). The second fiber is set aside as spare to protect the working fiber. Whenever a failure is detected, the system automatically switches the affected traffic from the working fiber onto the protection one as illustrated in Figure 1. For any communication between a pair of nodes (A,C), the signal from A to C travels on one side of the working fiber ring and the return signal form C to A travels in the same direction but on the opposite side of that ring. Hence, every traffic demand travels the whole circumference and the maximum link flow on the links of the ring is the sum of all demands passing through the ring. Therefore the capacity required for both the working and protection fibers is the sum of all demands carried by the ring. This type of ring is sometimes referred to as a *dedicated protection* SHR since for each demand there is a corresponding amount of spare capacity set aside (i.e. dedicated) specifically to protect it.

In the second type of SHR, known as *bidirectional* SHR, not all working demands need to be routed in the same direction. In a two fiber implementation of such a ring, the bandwidth of each fiber is divided into two equal parts: one serving as working capacity and the other as spare. Both fibers may therefore be used to carry working traffic. For instance, in Figure 2, both channels of the demand A-C travel on the two fibers along the upper side of the ring, while those of A-D travel along the lower side. Hence, each demand only uses up capacity on the links on which it is routed, leaving the capacity on the rest of the links available for other demands. The capacity is therefore shared among the different demands carried by the ring, which is why it is also referred to as *shared protection* SHR. In the case of a failure, the affected traffic is switched to the spare capacity as illustrated in Figure 2.

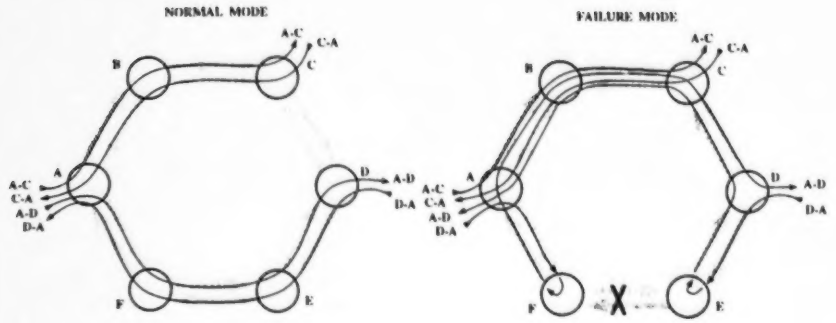

Figure 2: Bidirectional or Shared Protection SHR.

Because of the division into working and protection bandwidth, the overall capacity requirement for each fiber is two times the maximum flow passing through any of its links. By routing the demands along the ring so as to minimize this maximum flow, one may therefore achieve reductions in capacity requirements when compared with the unidirectional SHR. Indeed, if the traffic pattern is relatively uniform between the nodes of the ring then a bidirectional SHR will be more efficient capacity-wise. However, if the traffic is essentially *hubbed*, then there will be no advantage over the unidirectional SHR. This subproblem of *balancing* the flow among the edges of a bidirectional SHR is in itself a very difficult problem. For more details the reader is referred to the following: Cosares and Saniee (1994), Dell'Amico, Labbé and Maffioli [7], Myung, Kim and Tcha [20], and Schrijver, Seymour and Winkler [23].

Planning the introduction of self-healing rings in transmission networks is a complex matter that is being studied by a growing number of researchers. This planning problem has been approached from various angles giving rise to quite different design problems depending on the simplifying assumptions considered and what characteristics the resulting survivable architecture should have.

One of the simplest versions of the problem addresses a situation where the network is composed of a single cycle (or topological ring) passing through all the nodes and on which several stacked SHRs must be defined. The network contains a single designated hub node where all SHRs interconnect, allowing demands to flow from one ring to another. Each node can be connected to several SHRs, requiring the placement of an ADM each time. In this stacked ring network design problem, one needs to determine which nodes should be connected to each SHR and how to route the traffic through the rings in order to minimize costs (see Sutter, Vanderbeck and Wolsey [30], Armony *et al.* [3]).

A second class of design problems arises when one requires that the architecture of the resulting network contains multiple rings corresponding to several

different topological cycles. This class of problems is further divided depending on whether the different cycles or clusters are required to be disjoint with respect to the demand nodes they connect or not, the former being considered as the "simpler" version (see Klincewicz, Luss and Yan [16], Altinkemer [1], Altinkemer and Kim [2], Goldschmidt, Laugier and Olinick [12], Semal and Wirl [22], Sutter and Fullsack [29], Di Lascio, Gambaro and Mocci [8]).

In this paper, we consider the latter type of multiple SHR network design problems which deal with networks in which the node clusters defining rings are not required to be disjoint (i.e. nodes can be connected to more than one ring) and the resulting architecture is not hierarchical by definition (i.e. inter-ring traffic is routed on the SHRs defining the network). Some of the earlier and better known work on this topic was carried out at Bellcore and resulted in the planning software package known as SONET Toolkit, see Cosares *et al.* [6], Wasem, Wu and Cardwell [33], Wasem, Cardwell and Wu [32], Cosares and Saniee [5]. It is essentially a "greedy" approach that sequentially constructs a survivable network design. It is based on a three level decomposition of the overall planning problem that can be described as follows. In the first stage, clusters of nodes that would make "good" SHRs are identified. These clusters are made up of nodes that share a significant amount of traffic between them and are geographically close together (i.e. thus defining a *community of interests*). Then, a heuristic procedure determines if the network contains a reasonable cycle connecting the cluster nodes to form a SHR (Wasem [31]). If such a cycle exists then the "ring" is included in the final design, otherwise the cluster is rejected. While there are still demands to protect, the procedure tries to identify other interesting clusters and so on. If some demands are left unprotected after completing the ring selection phase, then the software tries to protect them by routing them on (up to three) interconnected rings or, finally, by assigning them to a diverse protection mechanism. Once all demands have been protected with such structure, the last step of the approach determines the required size for the different SHRs and eventual protection mechanisms that were selected.

Within this same class, Laguna [18] considers the clustering part of the problem that is grouping demand nodes to form individual SHRs. Each node may be connected to several rings. The resulting SHRs are only "logical" structures since no steps are taken in the procedure to guarantee that there exists a feasible cycle connecting the nodes of the ring. All demands must be protected either by being routed within a single ring or, when no ring contains both their end-nodes, by routing them on two different rings, each containing one of the end-nodes. No assumption or restriction is made regarding inter-ring traffic, but a linear cost is associated to it which makes this equivalent to considering that all traffic between any two SHRs pass through a designated hub node. Finally, only unidirectional SHRs (of different capacities) are considered. The objective is to find a set of

logical SHRs that optimizes the trade-offs between ADM placement and inter-ring traffic costs. The author proposes a tabu search heuristic to find feasible solutions.

In this paper, we consider the same design problem but in which inter-ring traffic has to be accounted for explicitly and routed on the SHR network being designed. As in Laguna [18], we also consider unidirectional SHRs of different sizes but in addition we include in our model *hop* constraints and *dual homing* constraints — i.e. constraints forcing each SHR to have at least two different interconnection nodes with other SHRs so that all inter-ring traffic is protected from intermediate node failures. We propose a tabu search algorithm to find good feasible solutions.

To complete this overview, note that contrary to the preceding contributions which were generally intended for application at the local or regional level and used mostly unidirectional rings, some papers deal with variants of the multiple SHR network design problem that are more related to backbone transmission network applications and which use exclusively bidirectional SHRs (see Slewinsky, Grover and MacGregor [24] , Kennington, Nair and Rahman [15], Luss, Rosenwein and Wong [19], Labbé, Laporte and Soriano [17], Pesant and Soriano [21]). We refer the reader to Soriano *et al.* [27] for a more detailed overview.

The remainder of this paper is organized as follows. In the next section, we present the specifics of the problem at hand and a mathematical model for the logical design. The solution approach is described in Section 3 and computational results are presented in Section 4. Finally, to simplify the presentation and given our european industrial partner, we will use only SDH terminology. However, everything translates directly into SONET context.

2 Model and solution approach

The Self-Healing Ring Network Design Problem (or $RNDP$ in short) can be stated as follows: given a set of demand nodes N, a two-connected (meshed) network $G = (N, E)$ connecting these nodes via a set of edges (fiber optic links) E, an O-D demand matrix $F = (f^k)_{k \in K}$, where $k = (i, j)$ and f^k is the traffic demand between the origin $i = o(k)$ and destination $j = d(k)$ nodes of commodity k, with $i, j \in N$, one wants to determine a set of feasible SHRs that protects all demands at minimal cost. Let us now characterize more precisely the $RNDP$ variant studied here.

- We consider only unidirectional SHRs but of several capacities. We denote

by R the set of possible rings and by u_r the capacity (i.e. size) of ring $r \in R$.

- Nodes can be connected to more than one ring. A node is connected to a ring $r \in R$ by the installation of an ADM. This induces a fixed cost c_r depending on the ring size.
- Hub constraints: all traffic between any two SHRs has to pass through a hub node (or interconnection node or zonal center) belonging to both SHRs. We denote by Z the subset of these nodes.
- A demand can be split on several paths (i.e. multirouting is allowed). Furthermore, there is no arbitrary upper limit on the number of rings used by any given demand path.
- Each SHR is required to contain at least two different interconnection nodes (dual homing constraints).
- There are no more than a pre-determined number Q of different nodes connected to any given ring (hop count limit).

In addition, the application considered here deals with transmission networks at the regional level. In such contexts, there are generally no really stringent restrictions with respect to fiber availability or ring layout within the cable network. Therefore, the physical realisation of a given "logical" SHR, as defined by the subset of demand nodes to be connected, is very often determined a posteriori, once the decision of which nodes to interconnect with that particular ring has been taken (see Cosares *et al.* [6], Laguna [18] and Luss *et al.* [19]). The following model also uses this simplifying assumption. We will therefore concentrate on the logical design (LD) aspect of the problem, that is determining which nodes should be connected by which SHR and how each individual demand flowing on the network is to be protected. The problem can then be formulated as a large integer program. Let $x_{ir} = 1$ if an ADM is placed at node i connecting it to ring r and 0 otherwise, $y_r = 1$ if ring r is used and 0 otherwise, v_{ir}^k (respectively v_{rj}^k) be the integer flow variables representing the amount of traffic of commodity k that accesses (respectively leaves) ring r at node $i = o(k)$ (respectively $j = d(k)$), and w_{rs}^{kz} the amount of inter-ring traffic of commodity k that passes from ring r to ring s at interconnection node z. The logical design of SHRs can be formulated as :

Formulation (LD)

$$z_{LD}^* = \min \sum_{r \in R} c_r \sum_{i \in N} x_{ir} + b \sum_{r \in R} \sum_{s \in R:\, s \neq r} \sum_{k \in K} \sum_{z \in Z} w_{rs}^{kz} \qquad (1)$$

Flow constraints

$$\sum_{r \in R} v_{ir}^k = f^k \quad k \in K,\ i = o(k) \tag{2}$$

$$\sum_{r \in R} v_{rj}^k = f^k \quad k \in K,\ j = d(k) \tag{3}$$

$$v_{ir}^k + \sum_{s \in R:\, s \neq r} \sum_{z \in Z \cap r \cap s} w_{sr}^{kz} = v_{rj}^k + \sum_{s \in R:\, s \neq r} \sum_{z \in Z \cap r \cap s} w_{rs}^{kz} \tag{4}$$

$$r \in R,\ k \in K :\ i = o(k),\ j = d(k)$$

Capacity constraints

$$\sum_{k \in K} (v_{o(k)r}^k + \sum_{s \in R:\, s \neq r} \sum_{z \in Z \cap r \cap s} w_{sr}^{kz}) \leq u_r \quad r \in R \tag{5}$$

Design constraints

$$\sum_{k \in K:\, o(k)=i} v_{ir}^k + \sum_{k \in K:\, d(k)=i} v_{ri}^k \leq \sum_{\substack{k \in K: \\ o(k)=i \text{ or } d(k)=i}} f^k\, x_{ir} \quad i \in N,\ r \in R \tag{6}$$

$$\sum_{i \in N} x_{ir} \geq 2\, y_r \quad r \in R \tag{7}$$

$$\sum_{i \in N} x_{ir} \leq Q\, y_r \quad r \in R \tag{8}$$

$$x_{ir},\ y_r \in \{0, 1\} \quad r \in R,\ i \in N \tag{9}$$

$$v_{ir}^k,\ v_{rj}^k \geq 0 \text{ integer} \quad r \in R,\ k \in K :\ i = o(k),\ j = d(k) \tag{10}$$

$$w_{rs}^{kz} \geq 0 \text{ integer} \quad r, s \in R :\ r \neq s,\ k \in K \text{ and } z \in Z \cap r \cap s \tag{11}$$

The objective (1) consists in minimizing the sum of ADM and inter-ring traffic costs. Constraints (2–4) are demand satisfaction and flow conservation constraints, constraints (5) are capacity constraints for the SHRs, and finally, constraints (6–8) are design constraints. Constraints (6) ensure that no demand accesses or leaves a given SHR at an origin or destination node if that node is not connected to that particular SHR, while (7) impose that each ring contains at least two interconnection nodes. Furthermore, constraints (8) ensure that there are no more than Q different nodes connected to any given ring (also termed

hop limit). Finally, it should be stressed that the different flow variables are constrained to take integer values because they represent high order multiplexed signals that cannot be split into non integer fractions for transmission. For instance, STM-4 signals can be split into several STM-1 but the STM-1 individual signals must stay whole.

For realistic sized instances, this model results in a huge integer program that is clearly not amenable to exact solution approaches. However, this logical design problem induces two interdependent decision problems: **(P1)** which deals with assigning demand nodes to the rings and **(P2)** which performs the routing of the resulting inter-ring traffic. Indeed, once the ADM connections are chosen, the problem reduces to a minimum cost multicommodity flow problem.

Our solution approach exploits this natural decomposition of the problem in an iterative procedure. The solution of P1 will define the data of P2, then the solution of P2 will allow us to modify the data of P1, and so on. We will therefore successively evaluate multiple ADM assignments to the rings by solving multicommodity flow problems. In this context, a tabu search algorithm that explores the space of possible assignments (x_{ir} variables) seems to hold good potential.

3 A Tabu Search heuristic

In this section, we present a tabu search heuristic developed to solve LD. We first detail the initialization step which builds a feasible solution. Then, we describe how inter-ring traffic is determined once the topology is fixed. We approximate the optimal routing cost by solving a multicommodity flow problem on some auxiliary graph. Finally, we present the neighbourhood structure, the solution attributes and the diversification mechanism used in our tabu search algorithm.

In the following, we will restrict ourselves to two ring sizes, namely STM-4 and STM-16. In the context of our industry partner, this was sufficient. However, considering more than two ring sizes (STM-64, etc.) does not present a challenge within our methodology, it only increases the overall size of the problem and therefore computing times. In these applications, larger equipments present economies of scale: STM-16 rings have 4 times the capacity of an STM-4 ring while STM-16 ADMs cost less than 2 times the cost of an STM-4 ADM. In addition, inter-ring traffic requires the installation of interfaces on both rings at the interconnection nodes, but the cost of such interfaces is much smaller than that of ADMs (i.e. roughly 25 times smaller than an STM-4 ADM).

3.1 Initialization step

The initialization procedure consists in a greedy heuristic that constructs a feasible solution of LD one ring at a time. For simplicity, we decided to build a solution without inter-ring traffic. However, to provide some free capacity in anticipation of the inter-ring traffic that will appear during the tabu search, we decided to limit the flow on a ring to a proportion α of its nominal capacity. In addition, since each ring contains at least two interconnection nodes, one is assured that demands between these nodes can be split over several rings, thus adding flexibility in view of the tabu search phase. The initialization step is outlined below.

Algorithm 1 (Initialization)

While there remain unassigned demands do

1. *Create a new ring and assign some unassigned demands to it.*
2. *Fill residual capacity of that ring by adding new nodes and demands while the used capacity stays below a proportion α of the capacity of the ring and the number of nodes does not exceed Q.*
3. *If the used capacity is below α times the capacity of the ring, add some unassigned demands between the interconnection nodes of the ring (if any) up to α times the capacity.*

End while

We now describe in detail the different steps of this algorithm.

1. *Creation of a new ring*

   We first determine the node having the maximum total demand for any pair of interconnection nodes. Then, we build a ring containing these three nodes and we assign it the corresponding demands. If there is no more unassigned demand between a node and an interconnection node, we create a ring with the origin and the destination of the highest unassigned demand, and two randomly chosen interconnection nodes. If the sum of demands not yet assigned is "large enough" (i.e. greater than a proportion β of an STM-16), the ring constructed is an STM-16, otherwise it is an STM-4.

   We denote by $dem(i, j)$ the demand between i and j not yet assigned to a ring and $demR(i; r)$ the total demand between i and nodes in ring r not yet

assigned to a ring. Given a ring r, $n(r)$ is the number of nodes connected to r and $usedcap(r)$ is the capacity of ring r already used for routing. The first step of the initialization phase can be summarized as follows.

- Let $z1, z2 \in Z$, and $i \in N \backslash Z$ be the nodes such that:

$$\begin{aligned} D^* &:= dem(i, z1) + dem(i, z2) \\ &:= max\{dem(a, z) + dem(a, z') : z, z' \in Z, a \in N \backslash Z\} \end{aligned}$$

- If $D^* \neq 0$,
  - create a ring r containing $z1, z2$ and i ($n(r) := 3$); if *Total unassigned demand* $\geq \beta$ *cap*(STM-16), r is an STM-16, otherwise it is an STM-4.
  - route demands $(i, z1)$ and $(i, z2)$ on r without exceeding the maximal proportion α.
- If $D^* = 0$,
  - let k be the highest demand not yet assigned,
  - choose randomly two interconnection nodes $z1$ and $z2$,
  - create a ring r containing $z1, z2, o(k)$ and $d(k)$ ($n(r) := 4$), and choose the capacity of r as above,
  - route demand k on r without exceeding maximal proportion α.
- Update $dem(.,.)$, $demR(.,.)$ and $usedcap(r)$.

2. *Filling capacity by adding new nodes*

Given the ring just created, we determine the node maximizing the total demand with nodes already in this ring, and we add it if we do not exceed the proportion α of capacity.

While ($usedcap(r) < \alpha\, u_r$ **and** $n(r) < Q$), **do**

- If $S := \{i : usedcap(r) + demR(i; r) \leq \alpha\, u_r\} \neq \emptyset$, then let $j := \arg\max\{demR(i; r) : i \in S\}$. Add j to r and assign the corresponding demands to r.
- Otherwise, fill the ring with interconnection node traffic as in Step 3. If no interconnection node traffic can be routed in r, if r is an STM-4, change it to an STM-16 and restart the whole procedure.
- If all the above fails, select the maximal unassigned demand k between two nodes not in r such that $usedcap(r) + f^k \leq \alpha\, u_r$, and assign it to r. If no such demand exists, start with a new ring.
- Update $dem(.,.)$, $demR(.,.)$, $n(r)$ and $usedcap(r)$.

End while

3. *Filling capacity with interconnection node traffic*

   If the previous operation does not fill the capacity up to a proportion α, and if there is some unassigned demand between the interconnection nodes that belong to the ring, we route this demand up to α times the capacity of the ring.

3.2 Routing and cost evaluation

Once subproblem P1 is solved, rings are determined and demand nodes are assigned to these rings. Subproblem P2 then consists in routing demands through these rings, verifying maximum ring capacity constraints and minimizing the inter-ring traffic cost. This problem reduces to a minimum cost multicommodity flow problem in an auxiliary graph presented in this section.

A demand between nodes i and j can be either routed on a ring containing both nodes i and j, or through 2 or more rings inducing then inter-ring traffic costs. In practice, this cost is a step function depending on the capacity of the interfaces required at the interconnection node but it can be reasonably approximated by a linear function with unit cost b.

The auxiliary graph $G^* = (N^*, A^*)$ is defined as follows. The set N^* contains 2 types of nodes:

- for each node $i \in N$, we define two nodes in N^* , the first one i will play the role of an "origin" point and the second one i' will correspond to a "destination";
- for each ring r, we define two nodes in N^* and we denote them by r and r'.

As illustrated in Figure 3, we consider 3 types of arcs.

- Node-ring arcs: (i, r) and (r', i') if node $i \in N$ is connected to ring r, which allow demands to access rings and exit from them at their end nodes.
- Ring arcs: (r, r'), for each ring, with a maximum capacity u_r (bold arcs in Figure 3). All the traffic routed on ring r will pass through this arc.
- Inter-ring arcs: (s', r) and (r', s) if rings r and s have a common zonal center (doted arcs in Figure 3), which allow inter-ring traffic to flow through the network.

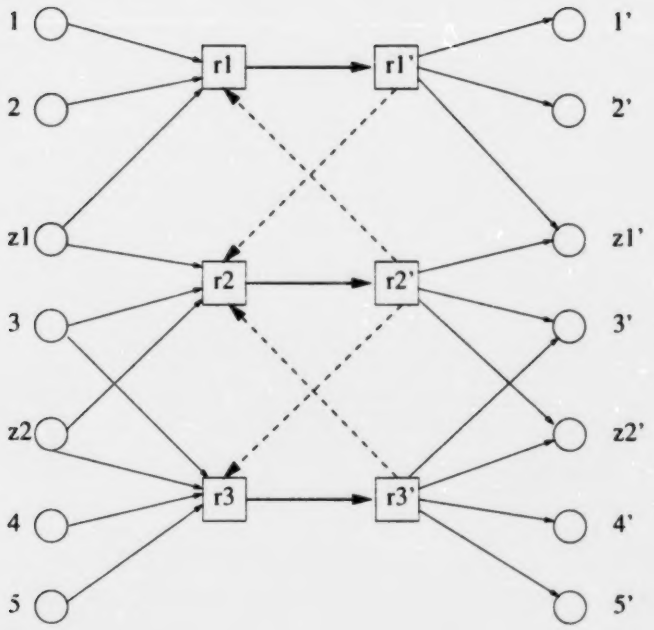

Figure 3: Auxiliary graph G^*.

Note that the maximum capacity of node-ring and inter-ring arcs is unlimited. For every traffic requirement between nodes i and j in N, we define an equivalent demand between nodes i and j' in N^*. Therefore, given an assignment of nodes to rings, subproblem P2 reduces to a minimum cost multicommodity flow problem in G^*, where the unit routing cost is zero for node-ring and ring arcs, while it is equal to b for inter-ring arcs. However, the existence of a feasible solution depends on the given node-ring assignment. Therefore, we add an infinite cost arc between each o-d pair with an unlimited maximum capacity. If such an arc appears in the solution, then the given node-ring assignment is not feasible.

3.3 Basic tabu search algorithm

First introduced by Glover [9] (and independently by Hansen [14]), tabu search is a general iterative search method that has been very successful in solving numerous types of hard combinatorial optimization problems (*cf.* Glover *et al.* [10], Glover and Laguna [11], and Soriano and Gendreau [26]). It is often referred to as a *metaheuristic* since it uses a set of rules to control an internal local search mechanism or neighbourhood search. The principle of the method is to continue the search for better solutions even if this produces a degradation of the objective function, thus preventing the method from getting trapped in local optima. Since

this can lead to cycling, tabu search uses the concept of "memory" to record the recent history of the search process and guide it away from previously visited solutions.

A solution S of the $RNDP$ problem is characterized by:

- a ring list $R(S) := \{r_1^S, ..., r_{n(S)}^S\}$,
  where r_j^S is the set of nodes connected to ring j by an ADM,
- a list of demand assignments to the rings $DA(S)$,
  to each demand corresponds a list of rings on which it is routed.

The total cost $c(S)$ of a solution is the sum of fixed cost $fc(S)$ (ADM installation) and inter-ring traffic cost $tc(S)$.

To define a neighbourhood $N(S)$ of such a solution S, we consider two types of modifications (or moves) of the ring list $R(S)$:

(1) remove node i from ring r_j^S (i.e. remove the ADM installed at i to connect it to r_j^S),

(2) add node i to ring r_j^S (i.e. place an ADM to connect i to r_j^S).

Moreover, when adding an ADM, we need to verify that we do not exceed the maximum number of ADMs per ring, and when suppressing the last node which is not an interconnection node, the ring is destroyed.

An optimal routing can then be obtained by solving a minimum cost multi-commodity flow problem on the auxiliary graph induced by the new node-ring assignment. However, this is a NP-hard problem; solving it at optimality would take too much time. Therefore, we only modify the routing of demands whose origin or destination is equal to i. The routing evaluation is then performed by solving a minimum cost transhipment problem with one origin and multiple destinations — one for each affected demand — in the auxiliary graph described in the previous section.

Hence, given a current solution S, the basic step of our tabu search algorithm consists in finding the best *feasible* neighbour, that is :

$$S^* := \arg\min\{c(S') : S' \in N(S)\}.$$

To prevent the procedure from cycling, the move, characterized by a pair (i, r_j^S), leading to the best neighbour is declared tabu for a number T of iterations. This

basic tabu search algorithm is outlined below. We denote by TL the tabu list, that is the set of node-ring pairs presently tabu.

Algorithm 2 (Basic Tabu Search)

1. *Initialization*

   - *Let S^0 be the feasible solution obtained with Algorithm 1.*
   - *Let $t := 0$, $S^* := S^0$ and $TL := \emptyset$.*

2. *Neighbourhood Evaluation*

   While the stopping criterion is not satisfied, do

   - *Evaluate neighbourhood $N(S^t)$.*
   - *Let S^{t+1} be the best feasible solution in $N(S^t)$ which is not tabu.*
   - *If $c(S^{t+1}) < c(S^*)$ then $S^* \leftarrow S^{t+1}$.*
   - *Update TL.*
   - $t \leftarrow t + 1$

   End while

We use two stopping criteria: a maximum number of iterations and a maximum computing time.

Initial experiments with this basic tabu search algorithm showed that complete evaluation of the neighbourhood was too time consuming. We therefore developed some criteria to reduce the size of the neighbourhood, to speed up its evaluation, as well as to improve the quality of solutions by introducing more flexibility.

- In order to use the residual capacity more efficiently when removing an ADM, it is often useful to reroute the interconnection node traffic as well as the demands affected by the ADM removal. To solve approximately this multicommodity flow problem, we use a greedy heuristic which sequentially solves a series of minimum cost transhipment problems corresponding to each node involved (origin or destination).

- As fixed costs are large compared to inter-ring traffic costs, removing an ADM often leads to a decrease in total cost, while adding one generally results in increased total cost. Hence, given a feasible solution S, we distinguish the neighbourhood $N_1(S)$ obtained by removing a node from a ring and the neighbourhood $N_2(S)$ obtained by adding a node to a ring.

Instead of evaluating the global neighbourhood $N(S) = N_1(S) \cup N_2(S)$, we first search neighbourhood N_1. If no feasible non-tabu solution is found, then neighbourhood N_2 is evaluated.

- To speed up the evaluation, we stop searching N_1 as soon as we obtain a feasible solution by removing an STM-16 without creating inter-ring traffic. Indeed, if such a neighbour exists it is certainly one of the best solutions in the neighbourhood. If no STM-16 can be removed, we stop the evaluation if a feasible solution is found by removing an STM-4 without creating inter-ring traffic.

- When searching neighbourhood N_1, if removing an ADM from a ring leads to an unfeasible solution, it is highly probable that this modification will stay unfeasible while moves are restricted to removals. Therefore, these moves are stored in a second unlimited tabu list until the end of this phase. When a new ADM is added, the list is emptied.

- During the second phase of the algorithm (neighbourhood N_2), the best choice for adding an ADM is always an STM-4 since it is more economical than an STM-16. Nevertheless, it turns out that this kind of move keeps the search in the same region of the solution space. To avoid this drawback and speed up the search, we perform a partial evaluation of N_2: for STM-16 all possible locations are tested while for STM-4 each node is added into only one randomly selected ring (instead of all possible rings).

3.4 Diversification phase

An important ingredient for local search efficiency is diversification. The aim of diversification is to guide the search process away from regions that have already been explored and into unexplored and hopefully promising areas of the solution space (see Glover and Laguna [11] and Soriano and Gendreau [25] for more details).

In our particular case, the basic search mechanism used by the tabu search algorithm just described cannot easily modify the number of rings present in a solution (i.e. it has to either empty a ring one ADM at a time or create a new ring having a single demand node and two interconnection nodes which is generally not a very attractive move costwise) and does not consider possible changes in the capacity of a given ring. The diversification procedure we propose enables the overall procedure to circumvent those limitations by evaluating the possible merger of two existing rings (possibly increasing the capacity of the resulting ring) or the splitting of an existing ring into two new rings (again possibly changing

the capacity). These types of more radical moves should in turn help the search process reach new regions of the solution space.

The main ideas that drive these modifications are the following:

1. If for a given ring r, a subset r_1 of its nodes of has only a small amount of traffic demand in common with the nodes in $r_2 := r \backslash r_1$, then splitting the ring into two smaller rings r_1 and r_2 may lead to a solution where the overall capacity required is smaller, therefore reducing the fixed cost without significantly increasing inter-ring traffic.

2. On the other hand, if two rings generate a large amount of inter-ring traffic, it could then be more economical to merge these rings into a single larger ring to prevent unnecessary inter-ring traffic to use up too much of the available capacity on the network.

The objective in splitting a ring in two is to minimize the inter-ring traffic generated. It is easy to see that this corresponds to solving a minimum-cut problem in a complete graph where the nodes correspond to demand nodes in the ring and each edge has a capacity equal to the demand between the nodes routed in the ring. However, it often turns out that the global minimum cut in this graph is given by a set containing only one node, while we would prefer to have a more balanced split. To avoid this drawback, we use the following heuristic: a minimum cut is computed for each pair of nodes using Gomory and Hu's algorithm [13], and we select the minimum cut corresponding to an edge not incident to a leaf in the Gomory-Hu tree. If no such cut exists however (meaning the Gomory-Hu tree is a star), the global minimum cut is chosen. We replace the ring by two new rings formed with the nodes on each side of the cut together with the interconnection nodes of the original ring. If it is feasible, we choose the two rings to be STM-4. Otherwise, we select the best combination of an STM-4 and an STM-16, and if none is feasible, both new rings are STM-16.

For merging, we select a pair of rings such that the sum of used capacities in both rings does not exceed an STM-16, the number of different nodes belonging to the pair does not exceed the threshold Q, and the interconnection nodes are the same in both rings. We then replace the original rings by a new one connecting all their nodes and we reroute on it all the corresponding demands and inter-ring flow. If the resulting flow is less than the capacity of an STM-4, the new ring is an STM-4, otherwise it is an STM-16.

The complete diversification procedure is summarized hereafter and can be seen as an additional last step within the *While* loop in Algorithm 2.

Algorithm 3 (Diversification procedure)

If best solution unimproved for** iter_div **iterations then do

1. *Try splitting each STM-16 ring in best solution and select best move, even if it results in a deterioration of solution quality (in order to have a minimal diversifying effect).*
2. *Repeat splitting step as long as it produces improvements (descent towards a local minimum with respect to the splitting neighbourhood).*
3. *Evaluate all ring pair mergers (as described above) and select best merge move. Implement the move if it improves the solution.*
4. *Repeat merge step as long as it produces improvements.*
5. *This final solution is set as the current solution and the basic tabu search resumes from it.*

4 Computational results

In this section we present and discuss the numerical experiments that were performed to calibrate and evaluate the TS algorithm presented in Section 3. We first describe the problem instances and then the specific experiments. All programs were implemented using the C++ computing language and were run on a PC with a Pentium III 466 Mhz processor running under Linux.

The set of test problems we considered contains *real* as well as *randomly* generated instances. A set of slightly "perturbed" real instances were provided to us by our industry partner. These 35 instances are grouped into 3 families according to the number of demand nodes ($n = 18, 32, 48$). All have exactly two interconnection nodes ($zc = 2$). They are identified in Table 1 by the indicator *real(n,zc)*.

In addition, we developed a generator that produces realistic random instances. Recall that in our context the physical realization of the rings is not considered to be restrictive, hence the data required to describe an instance simply consists in a symmetric demand requirement matrix. In real instances, there are usually three rather distinct levels for demand values: the largest values correspond to demands between interconnection nodes (i.e. zonal centers); the middle level corresponds to requirements between interconnection and *normal* demand nodes; the lowest demands are those between pairs of normal demand nodes with eventually several of them equal to zero.

To generate realistic random instances, we first specify the three parameters characterizing the network to be generated, that is the number of demand nodes n, the number of interconnection nodes zc, and the proportion p of all possible pairs of normal demand nodes for which there will be a non zero demand. Note that it is considered here that the only demands that can be zero are those between normal demand nodes. All other possible demands (i.e. between pairs of interconnection nodes and interconnection and normal demand nodes) are non zero. Then we associate to each node i an integer value $D_i \in [10, 15]$ if i is a zonal center, and $D_i \in [5, 10]$ otherwise. The demand between nodes i and j is then set equal to $D_i * D_j / r_{ij}$ rounded to the nearest integer and where r_{ij} is randomly generated between 2 and 10. We considered the different parameter combinations obtained with: n = 20, 30, 40; zc = 2 or 3; and p = 0.1, 0.3 and 0.5. For each combination, five different instances were generated thus providing us with 18 series of 5 instances each for a total of 90 random test problems.

The whole set of instances is summarized in Table 1. For each series of instances, we first give its identificator and then list the number of nodes, the number of zonal centers, the proportion of non zero demands of the third type, the number of instances and the problem type (*real* or *random*).

Series	n	zc	p	# instances	Type
$real(18, 2)$	18	2		11	real
$real(32, 2)$	32	2		18	real
$real(48, 2)$	48	2		6	real
$data(n, zc, p)$	20, 30, 40	2, 3	0.1, 0.3, 0.5	18 * 5	random

Table 1: Test problem characteristics.

A preliminary set of experiments was carried out to determine appropriate values for the different parameters of the procedure:

- the proportion threshold β which determines if a ring is an STM-16 or STM-4 in the initial solution,
- the proportion α which limits the flow on a ring in the initial solution,
- the tabu list size T,
- the maximum number *iter_div* of iterations without improvement before calling the diversification procedure.

The greedy heuristic (i.e. Algorithm 1) constructs a feasible solution without any inter-ring traffic. When considering the greedy heuristic in isolation, the best

results are of course obtained by setting α to 1 so that no portion of the ring capacity is set aside for possible inter-ring traffic (i.e. rings are filled as much as possible). As for the proportion threshold β controling the size of the ring being created, several values were tested. Low proportions (i.e. 0.3 or 0.5) tend to create more STM-16 rings, however these cannot be adequately filled afterwards with the simple constructive steps of the heuristic. Higher values of β (i.e. 0.7, 0.9 or 1) on the other hand, reduce the number of STM-16 rings and provide better capacity utilization in general. The best combination was clearly $\alpha = 1$ and $\beta = 0.9$.

In the context of the tabu search procedure, the initial solution needs to have some unused capacity in order to accomodate the inter-ring traffic that will be generated when exploring different alternate solutions. Setting $\beta = 0.7$ seemed to provide the required flexibility and this was then fixed for the rest of our experiments. We then tested several strategies for setting parameters α and T. Two stopping criteria were used simultaneously: a maximum number of iterations (1000) and a maximum computing time (1 hour) which was never reached. In fact, even the largest instances ran within 30 minutes of cpu time.

Results are summarized in Table 2. This table reports mean improvements, in percentage points, of solutions found by the basic tabu search procedure relative to the best solutions produced by the greedy heuristic. All values in a line are averages over all instances in the corresponding series. For example, the line $data(20)$ contains averages over all random instances with 20 nodes.

As can be seen from this table, the best results were obtained with parameters (T, α) set to (7 , 0.7). It appears that, given a fixed value of α, the procedure is rather robust with respect to the the tabu list size producing solutions of very similar quality for the different tabu list lengths. When analyzing the results in details, we observe that $T = 7$ is already sufficient to avoid cycling. Furthermore, for the same number of iterations, $T = 7$ allows to evaluate more solutions than $T = 15$, which explains why setting $T = 7$ or 10 gives slightly better results.

From Table 2, it is also clear that the value of α has an important impact on solution quality. For any given tabu list length, results are superior with high values of α, that is when rings in the initial solution are fuller. Clearly, the performance of the basic tabu search heuristic depends rather significantly on the initial solution. Moreover, for some small sized real instances, the solution it produces is inferior to the best solution found by the greedy heuristic.

As described in the previous section, the observation that the basic tabu search neighbourhood did not allow for easy changes in the number of rings composing a solution or to the capacity of a given ring motivated the introduction of a

(T, α)	(5 , 0.5)	(5 , 0.7)	(5 , 0.9)	(5 , 0.95)
$real(18, 2)$	-13,64	-15,91	-6,82	-4,55
$real(32, 2)$	-0,47	2,16	2,80	2,61
$real(48, 2)$	11,42	15,93	16,33	17,60
Real	-2,57	-1,16	2,10	2,93
$data(20)$	23 41	23,60	24,27	24,16
$data(30)$	30,68	32,51	32,21	31,36
$data(40)$	36,79	39,20	38,51	38,46
Random	30,29	31,77	31,66	31,33
Average	21,09	22,55	23,39	23,38
(T, α)	(7 , 0.5)	(7 , 0.7)	(7 , 0.9)	(7 , 0.95)
$real(18, 2)$	-13,64	-13,64	-4,55	-4,55
$real(32, 2)$	-1,00	1,26	2,67	2,53
$real(48, 2)$	13,68	14,79	16,43	17,87
Real	-2,45	-1,10	2,76	2,94
$data(20)$	22,34	22,73	24,17	23,82
$data(30)$	29,52	32,60	32,19	30,90
$data(40)$	36,74	38,22	38,55	38,62
Random	29,53	31,18	31,64	31,11
Average	20,58	22,14	23,55	23,22
(T, α)	(10 , 0.5)	(10 , 0.7)	(10 , 0.9)	(10 , 0.95)
$real(18, 2)$	-16,48	-11,36	-2,27	0,00
$real(32, 2)$	-0,95	1,53	2,73	2,35
$real(48, 2)$	10,09	14,54	15,88	17,69
Real	-3,94	-0,29	3,41	4,24
$data(20)$	21,41	22,21	22,86	23,76
$data(30)$	30,53	32,33	31,40	30,55
$data(40)$	37,49	38,68	38,19	38,20
Random	29,81	31,07	30,82	30,84
Average	20,36	22,29	23,14	23,39
(T, α)	(15 , 0.5)	(15 , 0.7)	(15 , 0.9)	(15 , 0.95)
$real(18, 2)$	-15,07	-13,64	-6,82	-4,55
$real(32, 2)$	-2,41	1,47	2,52	2,55
$real(48, 2)$	11,71	14,60	16,26	17,18
Real	-3,97	-1,03	1,94	2,83
$data(20)$	20,26	21,36	22,69	21,60
$data(30)$	30,30	31,53	30,84	29,87
$data(40)$	36,38	38,21	38,27	37,49
Random	28,98	30,37	30,60	29,65
Average	19,75	21,58	22,58	22,14

Table 2: Adjustment of parameters T and α (improvement %).

diversification procedure. To evaluate the impact of this procedure we carried out another set of experiments in which we fixed $T = 7$ (i.e. the best value according to the previous results) and considered two settings $\alpha = 0.7$ and 0.9. For each combination we then tested several values of *iter_div* (i.e. 30, 50 and 70 iterations without improvement). The results are summarized in Table 3 which again reports mean improvements, in percentage points, relative to the best greedy heuristic solutions. All values in a line are averages over all instances in the corresponding series. Note that to keep the table to a reasonable size, we do not report the individual results for each of the 18 $data(n, zc, p)$ series but present instead averages for each network parameter value over all series having that value in common. For example, the line $data(n, zc, 0.1)$ contains the averages over all random instances having a proportion $p = 0.1$ of non zero demands of the third type. The two columns under the heading *Tabu* reproduce the results obtained by the basic tabu search algorithm (i.e. without the diversification procedure) for $\alpha = 0.7$ and $\alpha = 0.9$ as they appear in Table 2.

	Tabu		Diversification					
α	0.7	0.9	0.7			0.9		
iter_div			30	50	70	30	50	70
$real(18, 2)$	-13.64	-4.55	0.00	0.00	0.00	0.00	0.00	0.00
$real(32, 2)$	1.26	2.67	7.95	6.87	7.65	7.56	7.29	7.60
$real(48, 2)$	14.79	16.43	11.83	11.68	12.61	12.23	12.26	12.30
Real	-1.10	2.76	6.11	5.54	6.09	5.98	5.85	6.02
$data(20)$	22.73	24.17	26.77	26.55	27.08	26.53	27.33	26.30
$data(30)$	32.60	32.19	35.90	35.91	35.94	36.09	35.65	35.92
$data(40)$	38.22	38.55	41.93	41.56	41.43	41.88	41.20	40.91
$data(n, 2)$	28.88	29.01	30.28	29.66	30.03	30.08	30.16	29.56
$data(n, 3)$	33.49	34.27	39.46	39.70	39.60	39.59	39.30	39.19
$data(n, zc, 0.1)$	24.85	26.89	17.47	17.22	17.59	17.66	17.47	17.07
$data(n, zc, 0.3)$	35.27	35.01	38.96	39.00	38.99	38.92	38.97	38.72
$data(n, zc, 0.5)$	33.43	33.02	48.16	47.81	47.87	47.92	47.75	47.33
Random	31.18	31.64	34.87	34.68	34.82	34.83	34.73	34.38
Average	22.14	23.55	26.82	26.52	26.77	26.76	26.64	26.43

Table 3: Solution quality for basic tabu search and diversification procedure for various parameter settings (improvement over greedy heuristic in %).

From Table 3, one can see that introducing the diversification mechanism significantly improves performance qualitywise. Even though for the series of larger real instances (i.e. $real(48, 2)$) and those of randomly generated instances with the sparsest demand requirement matrices (i.e. with $p = 0.1$) the diversification did not improve the average solution quality and even resulted in inferior solutions

for some instances, its overall impact is still clearly positive generating significant improvements in 90 out of 125 instances. Moreover, for the small sized real instances, it succeeds in matching the results obtaind by the best greedy heuristic, something the basic tabu search procedure could not accomplish on its own. The improvement in solution quality is more significant for random instances with 3 zonal centers and for those with the largest proportion of non zero demands ($p = 0.5$) which are the largest and more difficult instances in our test set.

As for the different settings for parameters α and *iter_div*, results are quite similar generating improvements in solution quality of about 26% on average over all instances. Again these results tend to confirm the robustness of the overall approach with respect to parameter settings. Nevertheless, it seems that calling diversification at a higher frequency leads to slightly better results.

Finally, as no lower bounding procedure is currently available for this problem, it is not possible to assess exactly the quality of the solutions found by the tabu search procedure. However, attaining over 26% improvement on average over a greedy constructive heuristic similar to the ones found in the industry is very encouraging and was deemed quite satisfactory by our industry partner.

5 Conclusion

In this paper, we propose a tabu search heuristic for the design of self-healing rings in a zonal network. The model studied comes from a real application with operational constraints such as dual homing and limited number of nodes per ring.

The algorithm is based on simple moves such as adding or removing a single node from a ring. Moreover, computational results presented on both real and randomly generated instances show the great benefit obtained by applying a diversification phase that modifies the best solution found using a more elaborate set of moves based on splitting rings in two and on merging existing rings.

This tabu search heuristic has demonstrated its efficiency by improving significantly upon the best solutions that could be found with a greedy constructive heuristic over a large set of real and randomly generated instances.

References

[1] Altinkemer, K. (1994). Topological Design of Ring Networks. *Computers and Operations Research*, 21:421–431.

[2] Altinkemer, K. and B. Kim. (1994). Heuristics for Ring Network Design. Working paper, Krannert Graduate School of Management, Purdue University.

[3] Armony, M., J.G. Klincewicz, H. Luss and M.B. Rosenwein. (1998). Design of Stacked Self-Healing Rings Using a Genetic Algorithm. Working paper AT&T Labs, Holmdel, NJ 07733, revised March 1998. To appear in *Journal of Heuristics*.

[4] Baudron, J., A. Khadr and F. Kocsis. (1993). Availability and Survivability of SDH Networks. *Alcatel Electrical Communications*, 4th Quarter, 339–348.

[5] Cosares, S. and I. Saniee. (1994). An Optimization Problem Related to Balancing Loads on SONET Rings. *Telecommunication Systems*, 3:165–181.

[6] Cosares, S., D.N. Deutsch, I. Saniee and O.J. Wasem. (1995). SONET Toolkit: A Decision Support System for Designing Robust and Cost-Effective Fiber-Optic Networks. *Interfaces*, 25:20–40.

[7] Dell'Amico, M., M. Labbé and F. Maffioli. (1996). Exact Solution of the Sonet Ring Loading Problem. Working Paper SMG 96/05, Service des Mathématiques de la Gestion, Université Libre de Bruxelles.

[8] Di Lascio, M., A. Gambaro and U. Mocci. (1994). Protection Strategies for SDH Networks. *Proceedings of the 6th International Network Planning Symposium - Networks'94*, Budapest, Hungary, September 4-9, pages 387–392.

[9] Glover, F. (1986). Future Paths for Integer Programming and Links to Artificial Intelligence. *Computers and Operations Research*, 5:533–549.

[10] Glover, F., E. Taillard, D. de Werra. (1993). A User's Guide to Tabu Search. *Annals of Operations Research*, 41:3–28.

[11] Glover, F., M. Laguna. (1997). *Tabu Search*. Kluwer, Norwell, MA.

[12] Goldschmidt, O., A. Laugier and E.V. Olinick. (1998). SONET/SDH Ring Assignment with Capacity Constraints. INFORMS Meeting, Montreal, Canada, April 26–29 (also working paper, Department of IE and OR, University of California, Berkeley).

[13] Gomory, R.E. and T.C Hu. (1961) Multi-terminal network flows. *SIAM Journal on Applied Mathematics*, 9:551–570.

[14] Hansen, P. (1986). The Steepest Ascent Mildest Descent Heuristic for Combinatorial Programming. *Congress on Numerical Methods in Combinatorial Optimization*, Capri, Italy.

[15] Kennington, J.L., V.S.S. Nair and M.H. Rahman. (1997). Optimization Based Algorithms for Finding Minimal Cost Ring Covers in Survivable Networks. Report 97-CSE-12, Department of Computer Science and Engineering, Southern Methodist University.

[16] Klincewicz, J.G., H. Luss and D.C.K. Yan. (1997). Designing Tributary Networks with Multiple Ring Families. Working paper AT&T Labs, Holmdel, NJ 07733, May 1997. To appear in *Computers and Operations Research*.

[17] Labbé, M., G. Laporte and P. Soriano. (1998). Covering a Graph with Cycles. *Computers and Operations Research*, 25:499–504.

[18] Laguna, M. (1994). Clustering for the Design of SONET Rings in Interoffice Telecommunications. *Management Science*, 40:1533–1541.

[19] Luss, H., M.B. Rosenwein and R.T. Wong. (1998). Topology Network Design for SONET Ring Architecture. To appear in *IEEE Transactions on Systems, Man and Cybernetics*.

[20] Myung, Y.-S., H.-G. Kim and D.-W. Tcha. (1997). Optimal Load Balancing on SONET Bidirectional Rings. *Operations Research*, 45:148–152.

[21] Pesant, G. and P. Soriano. (1998). An Optimal Strategy for the Constrained Cycle Cover Problem. Research report CRT 98-51, Centre de recherche sur les transports, Université de Montréal.

[22] Semal, P. and K. Wirl. (1994). Optimal Clustering and Ring Creation in the Network Planning System PHANET. *Proceedings of the 6th International Network Planning Symposium - Networks'94*. Budapest, Hungary, September 4-9, pages 303–308.

[23] Schrijver, A., P.D. Seymour and P. Winkler. (1998). The Ring Loading Problem. To appear in *SIAM Journal of Discrete Mathematics*.

[24] Slewinsky, J.B., W.D. Grover and M.H. MacGregor. (1993). An Algorithm for Survivable Network Design Employing Multiple Self-healing Rings. *Proceedings of IEEE GLOBECOM'93*, pages 1568–1573.

[25] Soriano, P., M. Gendreau. (1996). Diversification Strategies in Tabu Search Algorithms for the Maximum Clique Problem. *Annals of Operations Research*, 63:189–207.

[26] Soriano, P., M. Gendreau. (1997). Fondements et Applications des Méthodes de Recherche avec Tabous. *R.A.I.R.O.*, 31(2):133–159.

[27] Soriano, P., C. Wynants, S. Séguin, M. Labbé, M. Gendreau, B. Fortz. (1998). Design and Dimensioning of Survivable SDH Networks. In B. Sansò and P. Soriano (eds), *Telecommunications Network Planning*, Kluwer, Norwell, MA, pages 149–169.

[28] Sosnosky, J. and T.H. Wu. (1991). SONET Ring Applications for Survivable Fiber Loop Networks. *IEEE Communications Magazine*, June:51–58.

[29] Sutter, A. and J.L. Fullsack. (1994). SDH Network Planning in a Changing Environment. *Proceedings of the 6th International Network Planning Symposium - Networks'94*. Budapest, Hungary, September 4-9, pages 149–154.

[30] Sutter, A., F. Vanderbeck and L.A. Wolsey. (1998). Optimal Placement of Add/Drop Multiplexers: Heuristic and Exact Algorithms. To appear in *Operations Research*, 46(5).

[31] Wasem, O.J. (1991). An Algorithm for Designing Rings for Survivable Fiber Networks. *IEEE Transactions on Reliability*, 40:428–432.

[32] Wasem, O.J., R.H. Cardwell and T.H. Wu. (1992). Software for Designing Survivable Sonet Networks using Self-Healing Rings. *Proceedings of IEEE International Conference on Communication (ICC'92)*, pages 425–431.

[33] Wasem, O.J., T.H. Wu and R.H. Cardwell. (1994). Survivable SONET Networks - Design Methodologies. *IEEE Journal on Selected Areas in Communications*, 12:205–212.

[34] Wu, T.H., D.J. Kolar and R.H. Cardwell. (1988). Survivable Network Architectures for Broad-Band Fiber Optic Networks: Models and Performance Comparison. *IEEE Journal of Lightwave Technology*, 6:1698–1709.

[35] Wu, T.H. (1992). *Fiber Network Service Survivability*. Artech House, Boston, MA.